

Authentication Systems

This invention relates to authentication systems.

With the increasing usage of the internet as a business and social tool it is becoming more important that secure access to sensitive and personal information can be provided. Biometrics, the application of statistical analysis to identify an individual through their biological or physiological characteristics, is emerging as a key aspect in new security systems. Using biometrics it is possible to avoid pitfalls encountered by traditional security systems where users are required to keep a piece of information, such as a password, safe.

There are two types of biometric verification systems, classified by the type of biometric used; static or dynamic. Static biometric systems remain stable over time (barring injury), and examples of such biometric systems include fingerprinting systems, iris and retinal scan systems and hand geometry measurement systems. Dynamic biometric systems are subject to change over time, and examples of such biometric systems include signature systems, voice print systems and typing style systems. However such systems generally require specialised equipment at the point of use, thus rendering such systems unsuitable for multiple Internet based applications requiring secure authentication.

It is an object of the invention to provide an authentication system which is particularly suitable for multiple Internet based applications, as well as for a wide variety of other applications.

According to the present invention there is provided an authentication system for authenticating a user's signature as electronically inputted into the system by a manual input device providing an output indicative of its location with respect to time when manipulated by the user, the system comprising:

- (a) extraction means for extracting angle and distance data relating different parts of the user's signature inputted into the system by the input device;
- (b) reference means for storing corresponding angle and distance data relating to a reference signature;
- (c) comparison means for comparing the data extracted by the extraction means to the reference data stored by the reference means; and
- (d) verification means for providing an output indicative of an appropriate match between the inputted signature and the reference signature in dependence on the result of the comparison providing verification of the user's signature.

Such a system can provide an on-line dynamic biometric verification system that can be customised to multiple Internet based applications requiring secure authentication. The system requires no specialised equipment at the point of use, allowing access from any Internet capable computer with a mouse and Java compliant browser for example.

In this context it should be appreciated that the term "signature" is used in this specification to denote an electronic representation of an actual signature (the actual signature consisting of a distinctive representation of the user's name or any other distinctive pattern or representation, such as an emblem, mark or pictogram produced by the user), this electronic representation comprising in practice electronic data constituting an abstraction of the actual signature, for example by incorporating extracted angle and distance data relating to the signature as will be described more fully below. Furthermore the term "reference signature" is used to denote an electronic representation of a hypothetical authentic signature to which the inputted signature is to be compared, this hypothetical authentic signature comprising data constituting an abstraction of the actual signature extracted from a number of samples of the actual signature and possibly varying with time as further examples of the actual signature are sampled.

- (a) extraction means for extracting angle and distance data relating different parts of the user's signature inputted into the system by the input device;
- (b) registration means for setting up a reference data file compiled from angle and distance data relating to a plurality of samples of the user's signature inputted into the system by the user by means of the manual input device during a registration phase;
- (c) comparison means for comparing the angle and distance data extracted by the extraction means from the user's signature inputted into the system during an authentication phase to reference angle and distance data held in the reference data file, according to defined verification criteria ; and
- (d) verification means for providing an output indicative of an appropriate match between the inputted signature to be authenticated and the reference data in dependence on the result of the comparison, thereby providing verification of the user's signature.

Such a system can provide an on-line dynamic biometric verification system that can be customised to multiple Internet based applications requiring secure authentication. The system requires no specialised equipment at the point of use, allowing access from any Internet capable computer with a mouse and Java compliant browser for example.

In this context it should be appreciated that the term "signature" is used in this specification to denote an electronic representation of an actual signature (the actual signature consisting of a distinctive representation of the user's name or any other distinctive pattern or representation, such as an emblem, mark or pictogram produced by the user), this electronic representation comprising in practice electronic data constituting an abstraction of the actual signature, for example by incorporating extracted angle and distance data relating to the signature as will be described more fully below. Furthermore the term "reference signature" is used to denote an electronic representation of a hypothetical authentic signature to which the inputted signature is to be compared, this hypothetical authentic signature comprising data constituting an

2a

abstraction of the actual signature extracted from a number of samples of the actual signature and possibly varying with time as further examples of the actual signature are sampled.

It should also be appreciated that the term "authentication system" is to be interpreted as including within its scope not only systems for verifying a user's signature, for example for providing access to a bank account, but also systems for identifying individuals based on an input signature, for example in airport security. This identification is accomplished by comparing the input signature with stored reference signatures for a match above a degree of confidence.

In order that the invention may be more fully understood, reference will now be made, by way of example, to the accompanying drawings, in which:

Figure 1 is a diagram showing success steps in extraction of angle and distance data in relation to a user's signature in a system according to the invention;

Figure 2 is a diagram representing splitting of a user's signature for the purposes of data extraction;

Figures 3 and 4 are diagrams contrasting the data relationships obtained with the splitting of a user's signature in accordance with Figure 2 and according to a ranking approach;

Figure 5 is a flow diagram illustrating password and signature authentication in the system according to the invention;

Figure 6 is a block diagram of the system according to the invention;

Figures 7 and 8 are flow diagrams illustrating signature and password registration in the system according to the invention; and

Figure 9 is a flow diagram of training used in the system according to the invention.

The following description is given with reference to a preferred authentication system in accordance with the invention which has been shown in trials to give both a low false accept rate (FAR) and a low false reject rate (FRR). However it will be appreciated that many variations in such a system are possible within the scope of the invention, and that the choice of particular parameters, sample rates and verification procedures will depend on the particular application to which the system is to be applied. Furthermore the description of the system will be given with reference to accessing of the system over the Internet by a user making use of a mouse and keyboard connected to a personal computer (PC). Of course, other types of input device, such as a stylus tablet, can be used with systems in accordance with the invention, and such systems can be applied to any application in which signature authentication is required and are not simply limited to internet access applications.

The description of the system given below will be divided into a description of the manner in which the system authenticates a user, a description of the manner in which the user initially registers on the system, and a description of the manner in which the system is trained.

Authentication

A key feature of the preferred authentication system in accordance with the invention is the ability to match parameters of a signature inputted by the user into the system using the mouse with the corresponding parameters of a reference signature held within the system. Spatial co-ordinates are extracted by the system from the inputted user signature, as shown at A in Figure 1 for example, to obtain a signature trace as shown at B in Figure 1, each spatial co-ordinate being accompanied by a temporal value. This signature trace is then normalised such that it contains say 100 temporally equidistant points using linear time warping. The signature is normalized so that its arc length is 1 and so that the total time taken to produce the signature is 1.

From the normalised signature trace the system extracts say 10 relative angle parameters (an angle between vectors to two points) and say 10 relative distance parameters (a Euclidean distance between two points). In order to extract points that

provide a high between-class variance and low within-class variance (i.e. so that the features are as unique as possible) the system uses a genetic algorithm (GA) to extract these features from the normalized signature. The GA is evolved using standard crossover, mutation and selection, and can be defined using the following relationships.

The Euclidean distance, D_{ij} , between two points $S_i=(x_i, y_i)$ and $S_j=(x_j, y_j)$, where $[(0 \leq i < j < N) \vee (i=N-1, j=0)]$, in a linearly time warped signature containing N' points is given by:

$$(Eq. 1) \quad D_{ij} = \sqrt{(x_j - x_i)^2 + (y_j - y_i)^2}$$

An example of such a distance $D1$ between two points in the trace is shown at C in Figure 1.

The vector associated with a point, S_i , is obtained by the function $V(S_i)$, which returns either a vector from the previous point to the current point, or from the last point to the first point in the trace.

$$(Eq. 2) \quad \begin{aligned} V(S_i) &= \overline{S_i - S_{i-1}} & , \text{ if } i > 0 \\ V(S_i) &= \overline{S_i - S_{N'-1}} & , \text{ otherwise} \end{aligned}$$

where \overline{S} indicates vector normalisation. Two such vectors $V1$ and $V2$ are shown at C in Figure 1 by way of example.

The angle, A_{ij} , from S_i to S_j where $[(0 \leq i < j < N) \vee (i=N-1, j=0)]$, is obtained from the function ζ , which returns the clockwise angle between the two vectors (as shown at D in Figure 1 for the vectors $V1$ and $V2$):

$$(Eq. 3) \quad A_{ij} = \zeta(V(S_i), V(S_j))$$

where $V(S_x)$ is defined by Equation 2.

The fitness for a pair of points (i.e. a single gene) is given by:

$$(Eq. 4) \quad fitness(G) = \sqrt{\left(n \sum_{k=0}^n f(x_k)^2 - \left(\sum_{k=0}^n f(x_k) \right)^2 \right) / (n(n-1) - (j-i)/N)}$$

where G is a gene containing the reference number of two points in the signature, that is $G = (S_i, S_j)$ for signature point i and signature point j , and where n is

the number of signatures input by a specific user during registration and $f(x_k)$ is a function that calculates either the angle (A_{ij}) or distance (D_{ij}) between the points S_i and S_j .

The fitness for the set of features (i.e. the chromosome) is given by:

$$(Eq. 5) \quad CF = \frac{\sum_{i=1}^g fitness(G_i)}{g + \alpha + \beta + \chi}$$

where g is the number of genes in the population and $fitness(G_i)$ is defined by Equation 4. The function α penalises chromosome fitness in proportion to the fitness of the worst gene, using the function ($\text{Min}(1.0 - \text{standard deviation}(G_{0..i}))$). To ensure that relationships are well distributed the signature is divided into g sections, as shown in Figure 1 at B. The function β then returns a bonus for each section containing a *from* point (S_i in $G = (S_b, S_j)$) and the function χ returns a bonus for each section containing a *to* point (S_j in $G = (S_b, S_j)$).

The parameters that are extracted by the system from the signature are illustrated in Figure 1, and examples of relationships that are identified by these parameters are shown in Figures 2, 3 and 4. Although the GA usage is considered important it may be possible to extract the angle and distance relationships using other techniques. It is the obtaining of these angle and distance relationships such that they are sufficiently unique that is the most important criterion to obtain the required accuracy of authentication.

The set of angle values and the set of distance values extracted for a particular user are used to represent their signature. To this end the system incorporates two neural networks, each containing ten input nodes, that is one network for the 10 angle parameters and one network for the 10 distance parameters.

Referring to the flow diagram of Figure 5, in use of the system to authenticate a user's signature the user must enter their user name into the system by means of the keyboard to identify themselves to the system. The user's template file is retrieved by the system in response to entry of the user name, this template file having been previously encrypted using a standard encryption algorithm during compilation of the

template file as will be described below. The user must then enter their password. During this password entry the system records timing information in addition to the password string. Password verification is provided in a two-stage process, firstly by a string match step indicated at 2 and secondly by a keyboard dynamics (KD) verification step shown at 4 in Figure 5. In the event of a match between the inputted word and the password in the user's template file, logic 3 initiates the KD verification step at 4, whereas, in the event of such a match not being found, the logic 3 initiates a reject indication in a reject authentication step 6. Where a match is found the timing with which the password was entered by the user on the keyboard is then verified in the KD verification step shown at 4. In the event of a match between the inputted KD and the KD in the user's template file, logic 5 enables signature validation at 8, whereas, in the event of such a match not being found, the logic 5 initiates a reject indication in the reject authentication step 6. In this regard it should be noted that, even if a string match or KD match is not found, the user must still enter the signature. A reject message is then provided at the end of the input sequence so that an imposter is not able to determine what particular part of the input sequence has resulted in such rejection.

The user must then enter their signature by appropriately manipulating the mouse to initiate signature validation at 8 in Figure 5. Both spatial and temporal information in relation to inputting of the signature must be gathered for the system to function correctly, though this may not be the case if specialised hardware is used that provides equidistant timing samples (in which case the timing values will therefore be implicitly provided after normalisation). In the event of a match between the inputted signature trace and the reference signature trace in the user's template file, logic 9 initiates authenticity confirmation at 10, whereas, in the event of such a match not being found, the logic 9 initiates a reject indication in the reject authentication step 6.

Authenticity is verified provided that the user passes all three tests (string match, keyboard dynamics match and signature match). Below is a brief description of how the system verifies each of the tests. A simple string match is used to ensure that the password entered is the same as that supplied by the authentic user. From the user password input hold and latency time values and the total time are extracted. The hold

times represent the length of time each key is held down, and the latency times indicate the time from releasing of one key until pressing of the next key, with the total time being the time taken from pressing of the first key until releasing of the last key. Before being fed into a neural network the hold and latency times are normalised by the total time, that is each hold and latency value is divided by the total time to type the password. The input node size for the password neural network is therefore (*number of hold times + number of latency times*) and the actual inputs are the normalised hold and latency times. Each of the neural networks has a single output node that should output 0 if the user is identified as a forger and 1 if the user is authenticated.

It is feasible that different normalisation techniques could be used and more than one neural network used (one for hold and one for latency times for example). The important point is that the keyboard dynamics of the password input is used. Also different normalisation and pre-processing steps could be applied to the signature trace.

Registration

The neural network based system functions in three distinct modes, that is registration, training and authentication. During the registration phase new users are required to select a user name and input a chosen password and signature multiple times. The gathered biometric data is processed to extract salient information, with techniques including the use of a genetic algorithm. The details of the salient information used are then stored in a template file. During the training phase a novel technique is used to automatically generate forged samples. These forged samples, together with the authentic user samples, are provided to a back-propagation neural network, which is trained and stored upon the server. During authentication the user logs into the system via an applet that accepts a username, password and signature. The user template file, retrieved from the server, contains details of the salient features for the authentic user, which are then extracted from the input biometric data and sent to the server for verification. The data between client and server may be communicated safely because information is not transmitted from which a signature could be reconstructed.

Eavesdroppers may, therefore, intercept all transfers without compromising system integrity.

No two signatures are identical, even when signed by the same person. The lengths of the signature trace (in terms of the number of sampled points), the spatial size and temporal information will all vary. These differences are exaggerated by this system because input noise caused by variances in the provided sampling rate will distort the input signature data. The input signature trace therefore needs to be pre-processed to reduce the effect of these differences and to convert the trace into a standard format. Signature traces are pre-processed to normalise the arc length (signatures with disjoint segments are joined by the system to produce a single continuous arc). Next, the total time taken to produce the trace is normalised. Finally, the traces are linearly time warped to contain a pre-determined number of temporally equidistant points, typically 100, using the process described by L. Lee, "Neural Approaches for Human Signature Verification", Proc. 3rd International Conference Document Analysis and Machine Intelligence (TPAMI), vol. 15, No. 9. 1993, pp. 953-957.

It is possible to represent a signature using all information obtainable from the raw signature trace in a similar way to the keyboard dynamics' data. This is, however, undesirable because, due to the abundance of available information, much of the data will not provide a significant degree of uniqueness or consistency and the usage of such information could, therefore, prove to be counter-productive. Storing all of the information is also costly (in terms of space) and has implications for processing overheads when training networks and verifying signatures. Fortunately, it is possible to represent a signature by a number of extracted features rather than using all of the raw data. To this end the system uses an adaptation of a technique disclosed in Ozcan, E and Mohan, C (1996), "Shape Recognition Using Genetic Algorithms", Proceedings of the IEEE International Conference on Evolutionary Computation, Nagoya (Japan) May 1996, pp. 414-420, and Ozcan, E and Mohan, C (1998), "Steady State Memetic Algorithm For Partial Shape Matching", Proceedings of the IEEE 7th Annual Conference on Evolutionary Programming, March 1998, to perform partial spatial shape

matching, where relative angle and distance relationships between shape (signature) points are used. Equations 1 to 3 above define these relationships.

To use the extracted angle and distance information to characterise a signature trace a technique must be implemented to obtain both the salient angle and distance relationships from any input signature. To do this the technique must obtain an adequate set of points from the signature, from which relationships are extracted. This is performed with the intention of minimising within-class variance and maximising between-class variance, where within-class variance is the degree to which patterns belonging to the same class (user) differ and between-class variance is the degree to which patterns belonging to different classes differ.

If users are required to access the system upon an uncontrolled network, such as the Internet, then standard encryption techniques should be used to encode data transmissions. During most data transfers the data sent cannot be used to reconstruct a user signature as the only sufficiently unique relationships identified by the Genetic Algorithm (GA) are used. For example, only the values pertaining to the 10 angle and 10 distance parameters need be sent and it is not possible for the whole signature to be reconstructed from these parameters alone. The most dangerous time therefore is when registration data is being sent to a remote server prior to feature extraction (using the GA) and storage. This problem could be avoided either by providing a secure location for registration or by extracting the angle and distance parameters upon the local machine (although this may take some time).

During registration the user must choose a username that is available and appropriate. The user must then enter a password and signature multiple times so that the data may be used for the system to generalise. Figure 8 is a flow diagram indicating the validation and multiple inputting of the password, and the process by which the number of valid entries of the password is counted to arrive at a final count value after which the registration proceeds to the next stage. Figure 7 is a flow diagram indicating the validation and multiple inputting of the user's signature, and the process by which the number of valid entries of the signature is counted to arrive at a final count value

after which the registration proceeds to the next stage. A user template file is stored by the system containing the data gathered during the registration process. The spatial and temporal data obtained is used to train the neural networks for authentication purposes. Neural networks are used for authentication during the login phase.

Salient information to be used to identify a user is extracted from the data stored in the user's template file, and training data is created using the salient information extracted from the user's registration data in order to train the neural networks.

Training

Because the biometric information used in this system is liable to change with time (writing and typing styles change) the system must be able to adapt. This adaptation can be provided by performing periodic retraining of the entire system using data accumulated from successful logins. Alternatively data could be presented to the trained networks at each successful login, from which an output error is calculated. A single back-propagation pass is then performed, allowing gradual evolution of the networks.

Any neural network based verification system must contend with the issue of obtaining large amounts of training data needed to ensure a good ability to generalise. The authentic user may provide positive samples at the registration phase. However there are two main problems associated with obtaining false password and signature data from real people. Firstly the authentic user's password and signature must be made available to such people and secondly people willing to provide a sufficient number of good quality forged samples must be found. For these reasons a challenge for designing the system was to determine a technique for auto-generating a sufficient number of useful false data samples to allow effective neural network training to take place.

The biometric data obtained from user input can be considered as residing in a small sub-space of a much larger space. It is possible to authenticate an individual based upon whether their input data falls within this profile space region, but the

difficulty for any system is determining the appropriate size and shape of the authentic user's profile space. In this system the mean and standard deviation extracted from the data supplied at the time of registration (hold/latency times or angle/distance relationships) are used to provide an approximate model of the profile space. In each plane of dimensionality the mean values provide the centre of the profile space whilst the radii are based upon the standard deviation values. Other values such as the mean deviation could be used in the authentication process.

When using any verification technique it is evident that the most difficult forgeries to recognise are those that are very similar to authentic samples, lying close to the authentic user's profile space. Forgeries which reside further from the profile space can more easily be rejected by a verification system and therefore need fewer training sets with which the system can learn. To determine an optimal solution to the verification problem mainly false samples that lie close to the profile space boundaries within a boundary space region are used to train the system, with a few outlying samples to ensure correct modelling of the problem domain. The boundary space region is an enclosing sub-space whose radius is the same as the profile space radius (in each plane of dimensionality) and is set at a distance 0.25 times the radius.

To generate meaningful false samples for neural network training values must be generated that lie within the boundary space region. These false samples are generated using pseudo-random values for each axis, based upon the authentic user's characteristic patterns (additional true samples are generated within the profile space). Using this technique difficult forgeries are generated because they often lie outside the profile space in only one plane of dimensionality.

To perform verification the system uses three neural networks for each user trained using the back-propagation algorithm as disclosed in Bishop, C. "Neural Networks for Pattern Recognition", Oxford University Press, 1995. The first network uses hold and latency times to test typing style, and the second and third networks use angle and distance information to test the input signature. Separate angle and distance networks are preferably used because a combined network may be unable to correctly

model the problem domain. The networks are trained using the authentic user data input at the time of registration and automatically generated false and true samples using the technique described in the previous section.

When performing gradient descent on the networks it is possible to over-fit a problem such that the network remembers the input patterns rather than establishing an ability to generalise. The global minima of an error surface may provide a bad solution here because the input patterns are remembered. In order to combat this problem the system uses a validation set during training to test for an ability to generalise. Gradient descent is performed with respect to the training set but the previously unseen validation set is used to test for generalisation ability. To create the training, validation and testing sets the authentic user data (and auto-generated true data) is split between the three sets. False data is generated for each using boundary space generation, with the validation and testing sets using a boundary space slightly closer to the profile space than the training set so that performance and ability to generalise is assessed based upon more difficult samples.

Figure 9 is a flow chart showing the required training steps. Initially features are extracted from the keyboard dynamics of password input and further features are extracted from the user signature input. At 20 the centre in each dimension of the profile space region within which user true samples are expected to lie is identified, and at 21 the width in each dimension of the profile space region is identified. The region is based upon the mean and standard deviation values of the extracted salient data. This region could also be calculated using metrics other than mean and standard deviation.

At 22 the centre in each dimension of the boundary space region within which user true samples are expected to lie is identified, and at 23 the width in each dimension of the boundary space region is identified. The region is based upon the mean and standard deviation values of the extracted salient data. This region could also be calculated using metrics other than mean and standard deviation.

At 24 and 25 false and true training samples are generated within the boundary and profile space regions respectively and added to the authentic user data. Furthermore false samples are generated outside the boundary and profile regions at 26. This data is then used for training, validation and testing data. The proportion of each data type in each of the training sets may be altered (including the use of zero authentic samples in the testing set for example). At 27 a neural network is trained using the keyboard dynamics data. In this case the hold and latency times are used after normalisation, but the exact data used could be varied, i.e. hold times only or latency times only. Also separate neural networks could be trained for the hold times and for the latency times. At 28 neural networks are trained using the signature data. More particularly two neural networks are trained to verify angle and distance relationships within a signature. For authenticity to be confirmed both of these networks must confirm authenticity.

Figure 6 shows a preferred implementation of the template file 11 on the server consisting of neural networks 12, raw data 14, the password string 15 and the metadata 16. As previously indicated the neural networks 12 incorporate keyboard dynamics data 30 (hold/latency times), signature angle relationship data 31 and signature distance relationship data 32. Furthermore the metadata 16 incorporates points 33 from which the angle relationships are extracted and points 34 from which the distance relationships are extracted. The dotted box indicates that the raw data 14 will only be communicated to the server during registration and then will not be accessible. This data does not have to be communicated at all if the user registers at a secure location or if the GA extracts feature points on the client machine and then sends only the relevant data and metadata to the server. The other boxes show the data that may be communicated between the client and the server, although in an ideal situation only the metadata 16 needs to be sent to the client at login. The client would then extract the appropriate features and send these back to the server, with the result that the authentic signature could not be reconstructed from the transmitted messages.

It is a particular advantage of such a system that it uses a hybrid approach to verification, requiring an authenticity confirmation from both a typing style test and a signature match. The first stage of the system verifies authenticity based upon typing

style, with biometric information obtained from keyboard dynamics of the user's password input. Using the password input as the source of biometric data means the security benefits of standard password verification are enhanced, whilst no increase is placed upon the user's cognitive load. The second stage of the system verifies authenticity based upon an on-line signature match that uses temporal and spatial information. Unlike other systems which have been proposed in the past the system uses the mouse as the input device. Two existing user skills are therefore built upon; mouse use and signature writing. Although a pen-based system could be more desirable in terms of ease-of-use, this would mean that it would no longer be possible to access the system via the Internet using no specialised hardware.

Furthermore, by using passwords and signatures to gather biometric data, it is possible to avoid negative social stigmas, such as may be encountered, for example, in use of fingerprint systems. In a study conducted with 35 participants it was determined that 83% of people are happy to provide signatures as a means of verification, and of these 97% would be happy to provide their signature for use on the Internet.